arXiv:1611.03372v1 [cs.RO] 10 Nov 2016

# A stochastically verifiable autonomous control architecture with reasoning

Paolo Izzo, Hongyang Qu and Sandor M. Veres

Department of Automatic Control and Systems Engineering
University of Sheffield, Sheffield S1 3JD, UK
{pizzo1, h.qu, s.veres}@sheffield.ac.uk

November 11, 2016

**Abstract**

A new agent architecture called Limited Instruction Set Agent (LISA) is introduced for autonomous control. The new architecture is based on previous implementations of AgentSpeak and it is structurally simpler than its predecessors with the aim of facilitating design-time and run-time verification methods. The process of abstracting the LISA system to two different types of discrete probabilistic models (DTMC and MDP) is investigated and illustrated. The LISA system provides a tool for complete modelling of the agent and the environment for probabilistic verification. The agent program can be automatically compiled into a Discrete-Time Markov Chain (DTMC) or a Markov Decision Process (MDP) model for verification with Prism. The automatically generated Prism model can be used for both design-time and run-time verification. The run-time verification is investigated and illustrated in the LISA system as an internal modelling mechanism for prediction of future outcomes.

## 1 Introduction

Autonomous control is an area within control sciences that emerged by upgrading classical *feedback control* by decision making on what control references to use. The purpose of feedback control is to regulate a system in order to make it follow a predefined reference input. Autonomous controllers are designed to make decisions *what* reference signal to use and, more generally, what goals to achieve and *how* to achieve them. They do so by generating and executing plans of action that work toward goals [1, 2]. Autonomous controllers aim to introduce a certain level of "intelligence" in control systems, that is the ability of a system to act appropriately in an uncertain environment [3].

A first attempt towards autonomous decision-making software was initially made by using Object Oriented Programming (OOP). However the passive nature of objects in OOP, led to the development of active objects called "agents" [4], which implement decision-making processes. A formal description of autonomous agents can be found in [4–6].

One of the most widely used "anthropomorphic" approaches to the implementation of autonomous agents is the Belief-Desire-Intention (BDI) architecture [4,7]. BDI agent architectures are characterised by three large sets of atomic predicates: *Beliefs*, *Desires* and *Intentions*. The most known implementations of the BDI architecture are the Procedural Reasoning System (PRS) [8, 9] and *AgentSpeak* [10]. AgentSpeak fully embraces the philosophy of Agent Oriented Programming (AOP) [11], and it offers a customisable Java based interpreter.

Autonomous agents have a considerable potential for implementation in all sorts of different applications. However their introduction in real-world scenarios brings along safety concerns, creating the need for *model checking* [12]. An early attempt to BDI agent verification can be found in [13,14], where the authors present a translation software from *AgentSpeak* to either Promela or Java, and then use the associated model checkers Spin [15, 16] and Java PathFinder (JPF) [17]. A subsequent effort towards verifiable agents was made by Dennis *et al.* [18] with a BDI agent programming language called *Gwendolen*, which is implemented in the Agent Infrastructure Layer (AIL) [19, 20], a collection of Java classes intended for use in model checking agent programs, particularly with JPF. An evolution of JPF is Agent Java PathFinder (AJPF) [21], specifically designed to verify agent programs. However JPF and AJPF introduce a significant bottleneck in the workflow as the internal generation of the program model, which is created by executing all possible paths, is highly computationally expensive. In [22] it is suggested to alleviate this problem by using JPF to generate models of agent programs that can be executed in other model-checkers. This idea is further developed in [23], which shows how AJPF can be modified to output models in the input languages of Spin or Prism [24], a probabilistic model checker. All of the approaches towards agent verification to date, do not provide the user with a complete framework to build and verify a probabilistic model, and mostly they do not perform at a level suitable for real-time applications.

In this paper we introduce a new agent architecture called Limited Instruction Set Agent (LISA). The architecture of LISA is based on the three-layer architecture [25] and the agent program is an evolution of Jason [7,26]. The aim is to simplify the structure and the execution of the agent program in order to reduce the size of the state-space required to abstract it and ultimately allow for a fast verification process. The agent program is developed and described with *sEnglish* [27,28], a natural language programming interface. The use of sEnglish provides a way to define both the agent program

and the environment model in an intuitive, natural-language document. The document will then be automatically translated into PRISM source code for verification by probabilistic model checking. This is done by first proving that LISA can be abstracted as a DTMC or a MDP, based on design choices made by the user. We also propose the use of probabilistic model checking in PRISM to improve the non-deterministic decision making capabilities of the agent in a run-time verification process. Using run-time verification the agent is able to look into the consequences of its own choices by running model checking queries through the previously generated PRISM model.

# 2 Background

## 2.1 Rational Agents

An agent-based system is characterised by its *architecture*, a description of how the agent reasoning communicates with lower abstraction subsystems and ultimately with the environment. By analogy to previous definitions [4, 5, 27], we define the agent reasoning as follows.

**Definition 1** (**Rational agent**). *A* rational BDI agent *is defined as a tuple*

$$\mathcal{R} = \{\mathcal{F}, B, B_0, L, A, A_0, \Pi\}$$

*where:*

- $\mathcal{F} = \{p_1, p_2, \ldots, p_{n_p}\}$ *is the set of all predicates.*
- $B \subset \mathcal{F}$ *is the total atomic Beliefs set.*
- $B_0$ *is the Initial Beliefs set.*
- $L = \{l_1, l_2, \ldots\ l_{n_l}\}$ *is a set of logic-based implication rules on the predicates of* $B$*.*
- $A = \{a_1, a_2, \ldots, a_{n_a}\} \subset \mathcal{F} \setminus B$ *is a set of all available actions. Actions can be either* internal*, when they modify the Beliefs set to generate internal events, or* external*, when they are linked to external functions. Beliefs generated by internal actions are also called 'mental notes'.*
- $A_0$ *is the set of Initial Actions.*
- $\Pi = \{\pi_1, \pi_2, \ldots, \pi_{n_\pi}\}$ *is the set of executable plans or* plans library*. Each plan* $j$ *is a sequence* $\pi_j(\lambda_j)$*, with* $\lambda_j \in [0, n_{\lambda_j}]$ *being the* plan index*, where* $\pi(0)$ *is a logic statement called* triggering condition*, and* $\pi_j(\lambda_j)$ *with* $\lambda_j > 0$ *is an action from* $A$*.*

During an execution the agent also uses the following dynamic subsets of the sets defined above:

- $B[t] \subset B$ is the *Current Beliefs* set, the set of all beliefs available at time $t$. Beliefs can be negated with a '$\sim$' symbol.
- $E[t] \subset B$ is the *Current Events* set, which contains events available at time $t$. An *event* is a belief paired with either a '$+$' or a '$-$' operator to indicate that the belief is either added or removed.
- $D[t] \subset \Pi$ is the *Applicable Plans* or *Desires* set at time $t$, which contains plans triggered by current events.
- $I[t] \subset \Pi$ is the *Intentions* set, which contains plans that the agent is committed to execute.

The *triggering condition* of each plan of the plan library is composed by two parts: a *triggering event* and a *context*, a logic condition to be verified for the plan to apply. We write $B[t] \vDash c$ when the Current Beliefs set "satisfies" the expression $c$, or in other words when the conditions expressed by $c$ are true at time $t$. Note that in all our definitions and throughout the paper, time $t \in \mathbb{N}_{\geq 1}$ refers to the integer count of reasoning cycles.

Although different AOP languages implement the agent in different ways, generally speaking an agent program is iterative. Each iteration is called *reasoning cycle*. The reasoning cycle of the LISA system is explained in Sec. 3.

## 2.2 Model checking and verification

Probabilistic model checking is an automated verification method that aims to verify the correctness of probabilistic systems, by establishing if a desired property holds in a probabilistic model of the system [29]. For the purpose of this work we will consider models in particular: DTMC and MDP. Referring to [29–31] we give the following definitions:

**Definition 2** (Discrete-Time Markov Chain (DTMC))**.** *A (labelled) DTMC is a tuple $\mathcal{D} = (S, s_0, \boldsymbol{P}, L)$, where $S$ is a countable set of states, $s_0 \in S$ is the initial state, $\boldsymbol{P} : S \times S \to [0, 1]$ is a* Transition Probability Matrix *such that $\sum_{s' \in S} \boldsymbol{P}(s, s') = 1$ and $L : S \to \wp(\mathcal{F})$ is a labelling function that assigns to each state a set of atomic prepositions that are valid in the state.*

**Definition 3** (Markov Decision Process (MDP))**.** *A (labelled) MDP is a tuple $\mathcal{M} = (S, s_0, C, Step, L)$, where $S$ is a countable set of states, $s_0 \in S$ is the initial state, $C$ is an alphabet of choices with $C(s')$ being the set of choices available in any state $s'$, $Step : S \times C \to \mathrm{Dist}(S)$ is a probabilistic transition function with $\mathrm{Dist}(S)$ being the set of all probability distributions over $S$ and $L : S \to \wp(\mathcal{F})$ is a labelling function that assigns to each state a set of atomic prepositions that are valid in the state.*

Detailed explanation on the techniques used to perform model checking on probabilistic models goes beyond the scope of this paper. However we report here the syntax of the language used to write properties to verify with model checkers, which is called Probabilistic Computation Tree Logic (PCTL) [32].

**Definition 4** (Syntax of PCTL)**.**

$$\begin{aligned} \phi &::= \ \texttt{true} \mid a \mid \phi \wedge \phi \mid \neg\phi \mid \texttt{P}_{\bowtie p}[\psi] \\ \psi &::= \ \texttt{X}\ \phi \mid \phi\ \texttt{U}^{\leq k}\ \phi \end{aligned}$$

*where $a$ is an atomic proposition, $\bowtie \in \{\leq, <, >, \geq\}$ and $p \in [0, 1]$.*

We also allow the usual abbreviations such as '$\texttt{F}\phi$' (equivalent to '$\texttt{true}\texttt{U}\phi$'). A commonly used extension of PCTL is the addition of *quantitative* versions of the $\texttt{P}$ operator. For example $\texttt{P}_{=?}[\psi]$ asks: "what is the probability of $\psi$ holding?". In the same way we can add the operators $\texttt{P}_{\texttt{min}=?}[\psi]$ and $\texttt{P}_{\texttt{max}=?}[\psi]$ for MDP models that ask: "what is the minimum/maximum probability of $\psi$ holding?".

PCTL formulas can be extended with *reward* properties [31] by the addition of the *reward operator* $\texttt{R}_{\bowtie r}[\cdot]$ and the following state formulas:

$$\texttt{R}_{\bowtie r}[\texttt{C}^{\leq k}] \mid \texttt{R}_{\bowtie r}[\texttt{F}\ \phi] \tag{1}$$

where $\texttt{C}$ is the cumulative reward operator, $r \in \mathbb{R}$, $k \in \mathbb{N}$ and $\phi$ is a PCTL state formula.

# 3 The Limited Instruction Set Agent

The architecture of LISA, depicted in Fig. 1, is based on the three-layer architecture [25]. Each block with rounded corners is a collection of so called *skills* that the agent reasoning is able to execute when invoking actions. Note the hybrid nature of the system: the dotted lines represent *symbolic* flows of information, while the solid line represent *numeric* information.

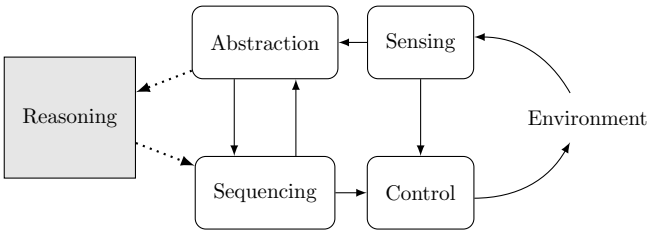


*Figure 1: The LISA architecture*

The agent program is an evolution of Jason [7, 26]. Here follows a brief overview of the modification that were made to Jason.

*Perception.* In LISA perception predicates can be of two types: *sensory perception* ($p \in B_s$) and *action feedbacks* ($p \in B_a$), therefore the Beliefs set is defined as:

$$B = \{B_s, B_a, B_m\} \tag{2}$$

where $B_m$ is the set of all possible mental notes. The *action feedbacks* are percepts that actions feedback to the Beliefs set of the agent in order to make the agent aware of the outcome of the action itself, i.e. success, partial success or failure. For the purpose of modelling, this classification is very important: the different nature of sensory percepts and action feedbacks needs to be modelled in a different way to accurately describe the behaviour of the environment. Messages are also handled as percepts.

*Goals.* In Jason there is a distinction between beliefs and goals. In a practical sense this distinction does not have a great influence: beliefs and goals can both trigger plans. For this reason in LISA we drop the definition of goals, by also defining goals as beliefs. This can simplify the process of generating a model directly from the agent code, by simplifying the syntax, and it also simplifies the modelling of the belief update process, by reducing the number of states required to describe it.

*Logic rules.* In Jason logic-based implication rules are present but yet not well implemented, to the point that the main text itself [7] advises against their use. In LISA we allow for rule to change the Beliefs set and therefore generate events. This feature potentially reduces the state space by allowing the definition of shorter plans with less actions.

In Fig. 2 we describe the reasoning cycle for LISA. The first step is to update the Current Beliefs set with the Beliefs Update Function ($f_{BU}$), based on percepts, messages and mental notes, where logic rules are also applied. The Belief Review Function ($f_{BR}$) then checks what changes have been made to the Current Beliefs set and it generates the new Events set. The function $f_P$ gathers all the plans from the Plan Library that are triggered by the current events, if the plan context is applicable to the Current Beliefs set, the plan is copied to the Desires set. An external function called *Plan Selection Function* ($F_O$) selects one plan for each event and it copies it from the Desires set to the Intentions set. Finally for every cycle the function $f_{act}$ executes the next action for each plan.

The general flow is similar to that of Jason with mainly one distinction: in every reasoning cycle the Jason agent only allows for the handling of a single event (selected with a function called *Event Selection Function* $F_E$), and then the execution of a single action from the Intentions set (selected with a function called *Intention Selection Function* $F_I$). In LISA we implement a multi-threaded work flow that allows the handling of multiple events, and then the execution of multiple actions at the same time. This implies that the *Desires* set becomes:

$$D[t] = \{D_1[t], \ldots, D_{n_e}[t]\} \tag{3}$$

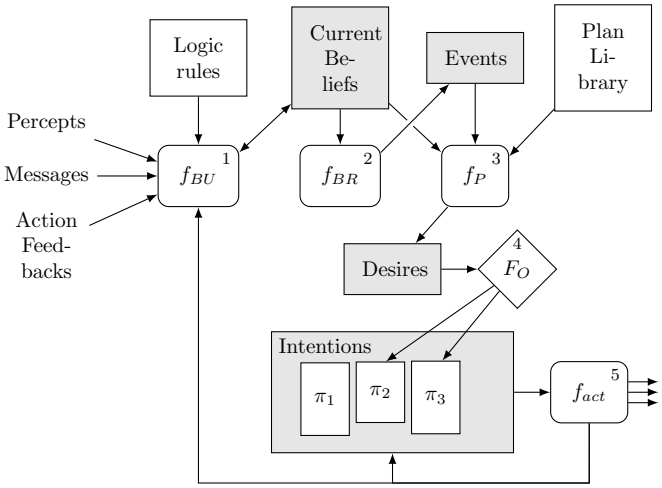


*Figure 2: The LISA reasoning cycle, rounded blocks represent internal functions, white square blocks are static sets, grey blocks are dynamic sets*

where each $D_j[t]$ is the set of plans triggered by an event $e_j \in E[t]$ and $n_e = |E[t]|$ is the number of events. Consequently, the function $F_O$, must be applied to every $D_j[t] \subset D[t]$. It is important to note that plans are *copied* into the Desires set from the Plan library, but not exclusively, which implies that different subsets of $D[t]$ may have a copy of the same plan. However if a plan is selected multiple times in the same reasoning cycle, it will only be executed once. Furthermore once a plan is selected from the Desires set and copied to the Intentions set for execution, if the plan is selected again in the future it will not be executed a second time, but it will carry on from the current state unless a plan interruption action is issued.

This multi-threaded implementation greatly simplifies the modelling process of the agent reasoning by drastically reducing the number of states required to describe it. By eliminating the need for specialised non-deterministic functions the model does not have to keep track of the events and actions activated in previous reasoning cycles therefore reducing the number of states. This also reduces the level of non-determinism in the agent reasoning, which then allows for a more precise generalisation of the abstraction process and in turn the application of an automatic modelling software that generates a complete and verifiable model directly from the agent code.

# 4 Abstraction to discrete finite-state machine

In this section we give a detailed description of the abstraction of the LISA reasoning to two kinds of discrete state machines: DTMC and MDP (see Definitions 2 and 3).

The agent defined in Definition 1 is in principle a deterministic system

with well defined rules and states. In the Jason implementation however, there are three functions that introduce non-determinism in the reasoning cycle ($F_E$, $F_O$ and $F_I$), which we reduce to one ($F_O$) with our LISA implementation. In Theorem 1 we show that the LISA system can still be modelled as a DTMC under the right conditions, and in Theorem 2 we show that the LISA system can always be modelled as a MDP.

In Definition 1 we introduced the concept of *plan* as a sequence $\pi = \{\pi(0), \pi(1), \ldots, \pi(n_\lambda)\}$. Assuming that a plan is not allowed to be executed multiple times in parallel, let us define a set of *plan indices* $\boldsymbol{\lambda}[t] = \{\lambda_1, \lambda_2, \ldots, \lambda_{n_\pi}\}$, which represents the state of all plans at time $t$. Note that, according to this definition, a plan $\pi_j$ is a member of the Intentions set $I[t]$ at time $t$ if and only if $\lambda_j > 0$ at time $t$. From $\boldsymbol{\lambda}[t]$ we can define a set of all possible indices as $\boldsymbol{\Lambda} = \{\boldsymbol{\Lambda}_1, \boldsymbol{\Lambda}_2, \ldots, \boldsymbol{\Lambda}_{n_\pi}\}$, where $\boldsymbol{\Lambda}_j = \{1, \ldots, n_{\lambda_j}\}$ is the set of natural numbers between 1 and the total number $n_{\lambda_j}$ of actions for each plan $\pi_j$.

**Theorem 1** (LISA abstraction to DTMC)**.** *Assuming the existence of sets of (discrete) probability distributions* $\mathrm{Dist}(B_s)$ *and* $\mathrm{Dist}(B_a)$*, over the set of percepts and the set of action feedbacks, if* $\forall\, i, j \in [1, n_\pi],\ \pi_i(0) \neq \pi_j(0)$ *the LISA can be modelled as a DTMC .*

*Proof.* A DTMC is completely characterised given a countable set of states $S$ and a transition function $\boldsymbol{P} :\ S \times S \to\ [0, 1]$. According to the definition of LISA, for a reasoning cycle to be completed the agent needs to be aware of $E[t]$ in order to recall plans from the plan library, $B[t]$ in order to check the plans context, and the state of the plans in $I[t]$ in order to execute the next actions. The state of a LISA is only relevant at the end of a reasoning cycle, therefore a generic state can be expressed as $s[t] = \{B[t], E[t], \boldsymbol{\lambda}[t]\}$. The state space, given by $S = \wp(B) \times \wp(B) \times \boldsymbol{\Lambda}$, is therefore finite and countable. The state of the agent is initialised by $s_0 = \{B_0, \emptyset, \mathbf{0}\}$ and by triggering the actions listed in the initial actions set $A_0$.
The transition function describes the way in which the state changes at every step. For each reasoning cycle, events can be generated from change in beliefs, namely *mental notes*, *action feedbacks* and *percepts*. Changes in mental notes are given by internal actions, which are known from the plan indices $\lambda$. Changes in action feedbacks and percept are given by known probability distributions. If $\forall\, i, j,\ \pi_i(0) \neq \pi_j(0)$, e.g. if all plans have different triggering conditions, then

$$\forall t \in \mathbb{N}_{\geq 1}\ ,\quad \left| \bigcup_{k=1}^{n_e} D_k[t] \right| = |D[t]| \leq |E[t]| \tag{4}$$

each event will trigger at most one plan, therefore $F_O$ becomes a trivial one-to-one mapping, therefore the system does not show any non-deterministic behaviour, hence the LISA reasoning can be modelled as a DTMC. □

**Theorem 2** (LISA abstraction to MDP)**.** *Assuming the existence of sets of (discrete) probability distributions* $\mathrm{Dist}(B_s)$ *and* $\mathrm{Dist}(B_a)$*, over the set of percepts and the set of action feedbacks, the any LISA reasoning can be modelled as a MDP.*

*Proof.* A MDP is completely described given a countable set of states $S$ and a transition function $Step : S \times C \to \mathrm{Dist}(S)$, with $C(s')$ being the set of available choices in any state $s'$. The set of states can be built as shown in Theorem 1.
If $\forall\, i, j \in [1, n_\pi],\ \pi_i(0) \neq \pi_j(0)$, according to Theorem 1, the system does not show any non-determinism. However, if $\exists i, j \in [1, n_\pi] : \pi_i(0) = \pi_j(0)$, then

$$\exists t' \in \mathbb{N}_{\geq 1} : \left| \bigcup_{k=1}^{n_e} D_k[t'] \right| > \left| E[t'] \right| \tag{5}$$

the number of applicable plans is greater than the number of events, therefore for some event $e_k[t']$ ($k \in [1, n_e]$), the application of the Plan Selection Function ($F_O(D_k[t']) = \pi$) involves a non-deterministic choice that implies different future probabilistic outcomes from action feedbacks, which prevents the modelling with DTMC. However this choice represents the only non-deterministic part of the agent, thus $C(s') = D[t']$. Once a choice is made by the Plan Selection Function, the transitions can be defined by changes in beliefs, given by internal actions and known probability distributions as shown in Theorem 1, and therefore the LISA modelling as a MDP is complete. □

Probabilistic models such as DTMCs and MDPs can be verified by means of probabilistic model checking, by using dedicated software such as Prism. Theorems 1 and 2 therefore imply that LISA can be verified, assuming that probability distributions of the percepts and action feedbacks are well defined. Theorems 1 and 2 also imply the availability of two options when designing the agent program: to design an agent with all unique triggering conditions, so to possibly improve model checking speed but requiring more effort from the designer, or design an agent with matching triggering conditions so to simplify the design but requiring more computation for the model checking.

## 5 Probabilistic modelling within agent programs

In this section we describe the process of modelling probabilistic behaviour of the environment and the action feedbacks in the agent code. The aim is to use a unified approach that allows to obtain a complete model of the agent and its interactions with the environment from a single document. The reasoning of the agent is implemented in sEnglish [27, 28], to which

we add a few features that give the programmer the option of defining the probabilistic parts of the system. Along with the probabilistic modelling we also introduce a *reward* structure which allows to define and use the reward properties supported by PRISM.

The action feedbacks are modelled within the action definition of sEnglish by defining the following three parameters: a *probability value* $p$, the *average number of reasoning cycles* $\mu$ in which the action feedback is expected to become true, and a *variance* $\sigma$. In this way we can simulate a time-delay-uncertain phenomenon without the need for real time models. For the percept process we use a similar notation with the possibility of defining probability distributions that are conditional to other beliefs. In particular the user defines: a *list of percepts or mental notes* to which the percept being modelled is conditioned to, probability, average number of reasoning cycles and variance of *activation* and *deactivation*.

The last feature we introduce is the possibility for the programmer to describe *reward* structures, that then allow to use reward properties as described in Equation 1. The reward values can be declared by adding a new '$\{\cdots\}$' structure to any percept declaration within the Percept Process section, or to any action within any of the executable plans.

By specifying all the necessary information, as described above, the designer is able to implement a complete model that includes a probabilistic description of the environment behaviour, e.g. percepts and action feedbacks. This allows to automatically generate PRISM input code for verification (see Sec. 6).

## 6 Design-time verification

The software used to perform the design-time and run-time verifications is PRISM [24, 33]. The modelling approach showed in Sec. 5 gives an sEnglish program that provides enough information to generate a complete PRISM model for verification. The PRISM model is generated here with a dedicated Matlab script. The translator only operates on the agent program itself, and it runs in the order of the tens of milliseconds on the laptop PC we used for the testing. For this reason the performances of the translator itself will be considered to be negligible for the results presented in this paper.

The automatically generated PRISM model is structured as follows: a variable is defined for every belief (percept, mental note and action feedback), a variable is also defined for every plan, representing the plan index $\lambda$ which captures the state of the plan at any given time. By using the synchronisation feature offered by the PRISM software, the reasoning cycle is simulated in two steps: a *Beliefs set update*, where variables associated with beliefs are updated, and a *plan index update*, where variables associated to plan indexes are updated according to the beliefs. With this method we en-

sure that plans only advance when the appropriate conditions on the Beliefs set are met. Note that there are no variables associated with actions as they are not part of the definition of state of the agent, as shown in Theorem 1.

Note that by using the approach presented in this paper, during the verification process, the user has access to every single belief and plan. This means that the property specification can touch any part of the system, allowing the user to define arbitrarily complex properties on any aspect of the reasoning process. This can be used to drastically reduce the design errors for autonomous agents.

For example, assume that an agent is implemented to have two opposite actions such as '`go left`' and '`go right`'. Assuming that the agent is programmed to have $\pi_2(1)$='`go left`' and $\pi_4(2)$='`go right`', the property:

$$\mathtt{P}_{\mathtt{max=?}}\ [\mathtt{F}\ \ (\mathtt{plan_2} = 1\ \&\ \mathtt{plan_4} = 2)]$$

will ask the model checker to generate "the maximum probability of '`go left`' and '`go right`' to be executed at the same time at some point in the future".

## 7 Run-time verification for improved decision-making

In this section we propose two different methods for using a run-time verification process as an internal model for improving the decision-making capabilities for the LISA system. The automatically generated PRISM, presented in Sec. 6, can also be used for run-time verification. Most of the computational power required to verify such a model is usually spent by the model checker when building the model itself, which does not influence the verification time. In other words, once the model is built, the user can run different verification queries without having to rebuild the model. In many cases, PRISM is able to compute the answers to those queries in a matter of seconds, even for a fairly complex model, therefore this can be a reasonable technique to use in this framework.

The first method is to implement the run-time verification process as a *skill* of the agent, e.g. as a module of the full system.The DTMC or MDP model is verified against a set of predefined queries. In particular, in PRISM, it is possible to check a query by selecting a starting state with the use of *filters* [33]. The run-time verification is then used to generate a set of results that will be interpreted by a '*generate beliefs*' function that will activate or deactivate certain beliefs in the agent Beliefs set.

The second method consists of implementing a *Plan Selection Function* that makes use of model checking to assess probability of success based on user-defined specifications, and selects the most suitable plan. A clear advantage to this approach is that, since the probabilistic model is generated

automatically, the user does not need to implement a specialised function for each agent. A possible implementation for this is as follows. The function takes as input the Current Beliefs and Desires set. The model generated at design time can be initialised with the current state and then checked against predefined queries. This results in probability values that can be used as indices to select the most likely to succeed plan amongst the ones in the Desires set.

Note that the two methods described here for run-time model checking, are not mutually exclusive: in case the programmer chooses to implement the LISA as a MDP, they could both be used at the same time.

## 8 A case study

Consider an Autonomous Surface Vehicle (ASV) designed for mine detection and disposal. The ASV is equipped with sensing equipment such as sonars and cameras that allow the detection of unidentified objects in the area of interest. These sensors give the vehicle a cone shaped visibility range. Using its pose in the environment and the information from the sensing equipment, the system is able to assess, on the fly, whether or not there has been any area left unclear. All the data collected is continuously sent back to the control centre. Once the mission is started, lower level tasks, such as for example collision avoidance, are carried out automatically by dedicated subsystems. During the exploration of the area, the system tags mine-like objects and logs their positions and available information for the human operators at the control centre to analyse and deliberate. In this scenario, a mission consists of a set of points (in terms of latitude and longitude) that outlines a specific area. An algorithm generates a sequence of waypoints connected by linear tracks, the parallel distance between tracks is calculated by considering the range of the available sensing equipment. We will call the linear tracks, and the area surrounding the tracks, "blocks".

In a best case scenario the exploration plan will be carried out as it is defined. However a number of problems can occur. In case that the weather condition becomes too harsh, the agent will wait for instructions from human operators. If the agent realises that there are areas left unexplored in the last block, it will make a non-deterministic decision on whether to immediately go back to re-explore missed spots, or keep going and come back at the end of the mission.

A fragment of the LISA program developed for this example is shown in Fig. 3. In Fig. 4 is shown a fragment of the Prism program that is automatically generated from the agent code. In Table 1 results are reported by running the model in Prism. All the testing was done by using an Apple laptop with a dual-core Intel Core i5-4258U 2.4GHz CPU and 16GB of memory running 64-bit Mac OS X 10.11.3. We implemented two different

```
PERCEPTION PROCESS
Monitor the following booleans:
//Percepts
Sea state is too high. {[],[0.5,10,0]}
I am at global waypoint.
Areas left unexplored.
Last waypoint reached. {[I am at global waypoint],[1,1,0]}
...
EXECUTABLE PLANS
...
//Plan 5
If ^[Block explored] while ^[Areas left unexplored] and ~^[
    Sea state is too high] then
  [Activate park mode.]
  [Generate set of waypoints.]
  +^[Re_exploring areas]
  [Activate drive mode.].
...
//Plan 8
If ^[Sea state is too high] while true then
  [Activate park mode.]
  [Wait for instructions.]
  +^[Waiting for instructions].
```

*Figure 3: Fragment of the agent program for this case study.*

```
module plan_5
plan_5: [0..4] init 0;
[t] plan_5=0 & !(plan_4=0 & block_explored=1 & (
    areas_left_unexplored=1 & sea_state_is_too_high=0)) -> (
    plan_5'=0);
[t] plan_5=0 & (plan_4=0 & block_explored=1 & (
    areas_left_unexplored=1 & sea_state_is_too_high=0)) -> (
    plan_5'=1);
//activate_park_mode
[t] plan_5=1 & !(park_mode=1) -> (plan_5'=1);
[t] plan_5=1 & (park_mode=1) -> (plan_5'=2);
//generate_set_of_waypoints
...
module wait_for_instructions
continue: [0..5] init 0;
abort: [0..1] init 0;
//continue[0.6,5,0] abort[0.4,5,0]
[p] !(plan_8=2) & (continue<=1 & abort<=1) -> (continue'=0) &
     (abort'=0);
[p] (plan_8=2) & (continue<=1 & abort<=1) -> (continue'=5);
```

*Figure 4: Fragment of the agent program for this case study.*

*Table 1: Verification results for the example. Both models are run with the following query:* $\mathtt{P}_{\mathtt{min}=?}\ [\mathtt{F} \leq 100\ \mathtt{mission_complete} = 1]$.

| Model | States | Transitions | Choices | Build time | Ver. time | Memory | Result |
|---|---|---|---|---|---|---|---|
| MDP | 270,268 | 420,431 | 276,454 | 38.061s | 1.901s | 10.0 MB | 0.6357 |
| DTMC | 157,072 | 231,148 | N/A | 38.146s | 2.052s | 3.5 MB | 0.6389 |

versions of agent programs for this case study: one by defining the agent to be abstracted as a DTMC, as per Theorem 1, and one as a MDP, as per Theorem 2. Both of the models feature 10 executable plans and 3 logic rules each, 4 percepts, 4 possible actions with a total of 5 action feedbacks, with additional conditions for the DTMC version.

Referring to Table 1, as expected, the MDP case generates a model that is larger than the DTMC counterpart. However even though the number of states in the MDP case is about 70% larger than the number of state for the DTMC model, the time required to build the model is very similar. This is possibly due to the way Prism handles the model building: the software constructs a MTBDD [34] structure, that is very much dependent on the logic structure of the model.

Both models were then ran with a standard verification query that calculates the minimum probability of completing the mission within 100 steps. Note that for the DTMC model the probability is a single value, which is still indicated here with 'minimum probability'. The verification time in this case is consistent with the increase in number of states from the DTMC model to the MDP model.

# 9 Conclusions

In this article we introduced a new architecture for BDI agent programming called Limited Instruction Set Agent (LISA). The architecture builds on previous implementations of AgentSpeak but with a simpler structure, in order to facilitate automatic verification of agent reasoning. The reasoning of the agent is defined in the Natural Language Programming (NLP) language *sEnglish*. We proved that the LISA architecture can be abstracted as a Markovian model (DTMC or MDP), we then showed how to define a full probabilistic model that includes the agent reasoning and the environment, all defined within the agent reasoning program. From the improved agent program we can then automatically generate a full probabilistic model in Prism's input language for verification.

The model generated from the agent code is used for both design-time and run-time verification. The design-time verification can serve to im-

prove and validate the agent design. The run-time verification can be used to improve the decision-making capabilities of the agent by implementing model-checking techniques as a means of simulation by the agent in order to predict future events and choose the most suitable strategy.

## Acknowledgment

This work was supported by the EPSRC project EP/J011894/2.

## References

[1] K. J. Aström, *Autonomous Control*. Berlin: Springer-Verlag, 1992, vol. Future Tendencies in Computer Science, Edts: A. Bensoussan and J.P. Verius, Vol. 653, Lecture Notes in Computer Science.

[2] K. J. Aström and R. M. Murray, *Feedback systems: an introduction for scientists and engineers*. Princeton university press, 2010.

[3] A. Meystel and J. Albus, *Intelligent Systems: architecture, design and control*. Wiley, 2002.

[4] S. M. Veres, L. Molnar, N. K. Lincoln, and C. Morice, "Autonomous vehicle control systems - a review of decision making," vol. 225, no. 2, pp. 155–195, 2011.

[5] M. Wooldridge, *An Introduction to MultiAgent Systems*. Chichester: Wiley, 2002.

[6] M. Wooldridge and N. R. Jennings, "Intelligent agents: theory and practice," *The Knowledge Engineering Review*, vol. 10, no. 2, pp. 115–152, 1995.

[7] R. H. Bordini, J. F. Hubner, and M. Wooldridge, *Programming multi-agent systems in AgentSpeak using Jason*. Chichester: Wiley, 2007.

[8] M. P. Georgeff and A. L. Lansky, "Procedural knowledge," *Proceedings of the IEEE*, vol. 74, no. 10, pp. 1383–1398, 1986.

[9] ——, "Reactive reasoning and planning." in *AAAI*, vol. 87, 1987, pp. 677–682.

[10] A. S. Rao, "Agentspeak (l): Bdi agents speak out in a logical computable language," in *Agents Breaking Away*. Springer, 1996, pp. 42–55.

[11] Y. Shoham, "Agent-oriented programming," *Artificial intelligence*, vol. 60, no. 1, pp. 51–92, 1993.

[12] E. M. Clarke, O. Grumberg, and D. Peled, *Model checking*. MIT press, 1999.

[13] R. H. Bordini, M. Fisher, C. Pardavila, and M. Wooldridge, "Model checking agentspeak," in *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*. ACM, 2003, pp. 409–416.

[14] R. H. Bordini, M. Fisher, W. Visser, and M. Wooldridge, "Verifying multi-agent programs by model checking," *Autonomous agents and multi-agent systems*, vol. 12, no. 2, pp. 239–256, 2006.

[15] G. J. Holzmann, *Design and Validation of Computer Protocols*. Prentice Hall, 1991.

[16] ——, "The model checker spin," *IEEE Transactions on software engineering*, vol. 23, no. 5, p. 279, 1997.

[17] W. Visser, K. Havelund, G. Brat, S. Park, and F. Lerda, "Model checking programs," *Automated Software Engineering*, vol. 10, no. 2, pp. 203–232, 2003.

[18] L. A. Dennis and B. Farwer, "Gwendolen: A bdi language for verifiable agents," in *Proceedings of the AISB 2008 Symposium on Logic and the Simulation of Interaction and Reasoning, Society for the Study of Artificial Intelligence and Simulation of Behaviour*, 2008, pp. 16–23.

[19] L. A. Dennis, B. Farwer, R. H. Bordini, and M. Fisher, "A flexible framework for verifying agent programs," in *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems-Volume 3*. International Foundation for Autonomous Agents and Multiagent Systems, 2008, pp. 1303–1306.

[20] L. A. Dennis and M. Fisher, "Programming verifiable heterogeneous agent systems," in *Programming Multi-Agent Systems*. Springer, 2008, pp. 40–55.

[21] L. A. Dennis, M. Fisher, M. P. Webster, and R. H. Bordini, "Model checking agent programming languages," *Automated software engineering*, vol. 19, no. 1, pp. 5–63, 2012.

[22] J. Hunter, F. Raimondi, N. Rungta, and R. Stocker, "A synergistic and extensible framework for multi-agent system verification," in *Proceedings of the 2013 international conference on Autonomous agents and multi-agent systems*. International Foundation for Autonomous Agents and Multiagent Systems, 2013, pp. 869–876.

[23] L. A. Dennis, M. Fisher, and M. Webster, "Two-stage agent program verification," *Journal of Logic and Computation*, p. exv003, 2015.

[24] M. Kwiatkowska, G. Norman, and D. Parker, "Prism 4.0: Verification of probabilistic real-time systems," in *Computer aided verification*. Springer, 2011, pp. 585–591.

[25] E. Gat, "On three-layer architectures," *Artificial Intelligence and Mobile Robots: Case Studies of Successful Robot Systems*, pp. 195–210, 1998, mIT Press.

[26] R. H. Bordini and J. F. Hubner, *Jason, A Java-based interpreter for an extended version of AgentSpeak*, 2007, manual version 0.9.5.

[27] N. K. Lincoln and S. M. Veres, "Natural language programming of complex robotic bdi agents," *Intelligent and Robotic Systems*, vol. 71, no. 2, pp. 211–230, 2013.

[28] S. M. Veres, *Natural Language Programming of Agents and Robotic Devices*. London: Sysbrain Ltd, 2008.

[29] M. Kwiatkowska and D. Parker, "Advances in probabilistic model checking," in *Verification, Model Checking and Abstract Interpretation (VMCAI)*. Springer, 2010.

[30] V. Forejt, M. Kwiatkowska, G. Norman, and D. Parker, "Automated verification techniques for probabilistic systems," in *Formal Methods for Eternal Networked Software Systems*. Springer, 2011, pp. 53–113.

[31] M. Kwiatkowska, G. Norman, and D. Parker, "Stochastic model checking," in *Formal methods for performance evaluation*. Springer, 2007, pp. 220–270.

[32] H. Hansson and B. Jonsson, "A logic for reasoning about time and reliability," *Formal aspects of computing*, vol. 6, no. 5, pp. 512–535, 1994.

[33] "Prism - probabilistic symbolic model checker." [Online]. Available: http://www.prismmodelchecker.org/

[34] M. Fujita, P. C. McGeer, and J.-Y. Yang, "Multi-terminal binary decision diagrams: An efficient data structure for matrix representation," *Formal methods in system design*, vol. 10, no. 2-3, pp. 149–169, 1997.